(22) Date de dépôt/Filing Date: 2000/04/25

(41) Mise à la disp. pub./Open to Public Insp.: 2001/10/21

(30) Priorité/Priority: 2000/04/21 (09/553,959) US

(51) Cl.Int.⁷/Int.Cl.⁷ G06F 17/60, G06F 17/30, G06F 3/033

(71) Demandeur/Applicant:
MALIBU ENGINEERING AND SOFTWARE LTD., CA

(72) Inventeurs/Inventors:
ANNAND, RITCHIE IAN, CA;
PERRY, GARRY ERWIN, CA

(74) Agent: SIM & MCBURNEY

(54) Titre : SYSTEME ET METHODE DE GESTION D'UNE BASE DE DONNEES A BASE DE DOCUMENTS
(54) Title: DOCUMENT-BASED DATABASE MANAGEMENT SYSTEM AND METHOD

(57) Abrégé/Abstract:
A method for managing data stored in a relational database involves developing a hierarchical model of a business schema. The model is then translated into hierarchical database tables for the database. The database tables model the business schema. Documents, including data controls having fields into which data is to be entered, are created using the hierarchy of the database tables. The data controls are live-linked to the database tables so that data entered into the fields automatically populates the database tables.

ABSTRACT

A method for managing data stored in a relational database involves developing a hierarchical model of a business schema. The model is then translated into hierarchical database tables for the database. The database tables model the business schema. Documents, including data controls having fields into which data is to be entered, are created using the hierarchy of the database tables. The data controls are live-linked to the database tables so that data entered into the fields automatically populates the database tables.

5

DOCUMENT-BASED DATABASE MANAGEMENT SYSTEM AND METHOD

Field Of The Invention

The present invention relates to data management and in particular to a document-based database management system and method.

5   Background Of The Invention

In many industries, companies include a number of corporate divisions that conduct business and report to a central headquarters. These companies tend to be large and accumulate significant amounts of data relating to their operations. In order for these companies to be run effectively, these companies must be able to use
10   the accumulated information to make educated business decisions.

For example, in the petroleum industry, oil companies must cope with vast amounts of data ranging from real-time data capture at a drilling rig, to pressure transient information, wellbore diagrams and facility inventories. Often, this information is stored in various locations and in a variety of databases. As a result,
15   much of this information is lost within the company and is not used during performance analysis at various levels throughout the company. This wealth of information therefore fails to impart knowledge and hence, intelligence to the company decision-makers.

To assist companies to deal with significant amounts of data, database
20   management systems have been developed. Unfortunately, many database management systems are written for vertical markets and provide only limited cross-talk between software applications accessing the databases. Also, these database management systems typically include embedded program code specific to the software applications that interface with them. This, of course, causes maintenance
25   problems whenever these software applications are updated, since it requires the program code of the database management systems to be modified.

Other systems have been designed to facilitate access to information stored in databases. For example, U.S. Patent No. 5,555,403 to Cambot et al discloses a data representation and query system that allows end-users to access and query a
30   relational database without knowing the relational structure of the SQL language. Data in the database is presented to the user in a form familiar to the business of the end-user through business objects. The system includes a query engine that generates

SQL statements in response to user input to query the relational database kernel of interest. The query engine is able to generate the SELECT, WHERE, FROM, HAVING, ORDER BY and GROUP BY clauses of an SQL sentence and to generate lists of joins involved in the query.

5   U.S. Patent No. 5,794,246 to Sankaran et al discloses a method and system for incremental aggregation of dynamically increasing database data sets. The system aggregates data stored within an operational database corresponding to newly received information to provide current information. Initially, an intermediate file is created that stores an aggregate data set including data values and count values that

10   correspond to specific group identifiers. Group identifiers within new sets of input data are compared with the group identifiers stored in the intermediate files. If a match occurs, the inputted data values are aggregated with the stored data value for that group identifier and the count value is incremented. If a match does not occur, the group identifier is added to the intermediate file and a count value of one is

15   appended to the added group identifier. Once all of the group identifiers in the new sets of input data are compared, all of the changes made to the intermediate file are stored in the aggregate data set.

U.S. Patent No. 6,038,558 to Powers et al discloses an architecture for extending data requesting and viewing capabilities of a database system. Data from

20   the database is requested and viewed using a plan. The plan defines the steps that need to be performed in order to request and view data. Each step includes executable code for accessing, manipulating and/or outputting data and conforms to a predefined interface that allows the step to be used in a plan, to manipulate data received from another step and to output data to another step.

25   Although the above references describe systems that facilitate access to data stored in databases, enhanced database management systems for managing data in a manner that conforms to the operations of businesses and simplifies use for end-users are continually being sought. It is therefore an object of the present invention to provide a novel document-based database management system and method.

30

Summary Of The Invention

According to one aspect of the present invention there is provided a document-based data management method comprising the steps of:

developing a model representing a business entity's schema for which

5　data is to be managed; and

creating documents using the model and live-linking the documents to a database structure corresponding to said model so that data entered into the documents populates the database structure.

In the preferred embodiment, the method also includes the step of

10　translating the model into the database structure for a database of interest. During the developing step, the model is developed using entity-relationship diagrams. During the translating step, database tables corresponding to the entity-relationship diagrams are generated. Documents are live-linked to the database structure through document objects that automatically populate the database structure with data entered into the

15　documents.

According to another aspect of the present invention there is provided a workbook comprising:

a user interface including a document selection panel listing at least one selectable document; and a document panel presenting a document selected from

20　said document selection panel, said document including at least one data control having a field to receive data; and

a linking engine live-linking said at least one data control to an underlying database, said linking engine automatically populating said database with data entered into the field of said at least one data control.

25　In the preferred embodiment, the linking engine is object-based and live-links the at least one data control to the database. The document selection panel includes a plurality of selectable documents that are held in folders listed in the document selection panel. Each folder holds user-specified documents. The database stores data hierarchically and the user interface includes a drill-down control to allow

30　information associated with different branches in the database at a selected level to be displayed within a document presented in the document panel.

According to yet another aspect of the present invention there is provided a modelling tool for creating database tables compliant with a backend database comprising:

a user interface including a model creation palette and user-selectable
5 controls to enable a user to create a hierarchical visual model representing a business entity, said model including entities within said business entity, attributes assigned to said entities and relationships between said entities; and

a table generator to create database tables for said backend database, said database tables following the hierarchy of said visual model.

10 In the preferred embodiment, the table generator creates a table object for each of the entities in the model. The table generator adds database field definitions to each of the table objects corresponding to the attributes assigned to the associated entity and corresponding to the relationship the associated entity has with another entity. The table objects are then used to create the database tables.

15 According to still yet another aspect of the present invention there is provided a document creation tool for creating a document including at least one data control having a field to receive data, said document creation tool comprising:

a user interface including a document creation palette presenting a document and user-selectable controls to enable a user to select data control types and
20 associated attributes from a data schema corresponding to a hierarchical model of a business entity and to create data controls on said document; and

a linking engine creating objects to live-link data controls created on said document to database tables of an underlying database corresponding to selected attributes of said data schema.

25 In the preferred embodiment, the linking engine automatically creates the objects when a selected attribute is placed on the document and an associated data control has been created. Attributes are placed on the document by dragging the attributes from a window and dropping the attributes on the document using a computer pointing device.

30 According to still yet another aspect of the present invention there is provided a method of licensing database management software to a licensee, said database management software including documents having data controls associated

-5-

with attributes assigned to entities within a model representing the business of said licensee, data to be managed being entered into a database through said data controls, said method comprising the step of:

charging a fee to said licensee for said software based on the number

5   of attributes assigned to entities within said model.

According to still yet another aspect of the present invention there is provided a method of licensing database management software to a licensee, said document-based database management system including documents created for said licensee, each document including at least one data control having an associated field

10   into which data is entered, said at least one data control being linked to a database and populating said database with data entered into said field, said data controls being associated with attributes assigned to entities within a model of the business of said licensee, said method comprising the step of:

charging a fee to said licensee based on the number of attributes

15   assigned to entities within said model.

According to still yet another aspect of the present invention there is provided a method for managing data stored in a relational database comprising the step of:

developing a hierarchical model of a business schema;

20   translating the model into hierarchical database tables for said database, said database tables modelling said business schema;

creating documents including data controls having fields into which data is to be entered using the hierarchy of said database tables; and

live-linking the data controls in said documents to the database tables

25   so that data entered into said fields automatically populates the database tables.

According to still yet another aspect of the present invention there is provided a method of transmitting data to a remote site over a telecommunications link comprising the steps of:

breaking a data file to be transmitted into packets;

30   transmitting each packet to said remote site, and for each transmitted packet, generating an entry corresponding thereto;

monitoring the amount of data in said file that has been transmitted to determine a current transmission point within said data file:

when said remote site acknowledges receipt of a transmitted packet, deleting the entry corresponding to said packet; and

5   if said telecommunications link terminates before said data file is fully transmitted, re-establishing said telecommunications link, re-transmitting packets associated with remaining entries and then resuming transmission of packets beginning from the current transmission point at the time the telecommunications link was terminated.

10   According to still yet another aspect of the present invention there is provided a file transmission application for transmitting data to a remote site over a telecommunications link comprising:

communications to establish a telecommunications link with said remote site;

15   a packetizer for breaking a data file to be transmitted into packets; and

a packet handler for transmitting each packet to said remote site, said packet handler generating an entry corresponding to each transmitted packet and monitoring the amount of data in said file that has been transmitted to determine a current transmission point within said data file, said packet handler further being

20 responsive to packet acknowledgments received from said remote site confirming receipt of packets and deleting entries corresponding to acknowledged packets, said packet handler re-transmitting packets associated with remaining entries and resuming transmission of packets beginning from the current transmission point when the communications re-establishes a telecommunications link that was terminated before

25 the data file was fully transmitted.

The present invention provides advantages in that data is managed across the complete spectrum of the business entity that models the way the end-user performs their business function as a part of the overall workflow of the business entity. This of course simplifies use of the database for the end-user since

30 information entered into documents automatically populates the database. Since the documents are created for each end-user, each end-user is exposed only to information of interest to them. Since data entered into documents automatically

populates the database, there is no confusion of where data is entered and reported. If the information does not show up in a given document, it is not in the database.

Brief Description Of The Drawings

An embodiment of the present invention will now be described more fully with reference to the accompanying drawings in which:

Figure 1 is a schematic diagram of a document-based database management system in accordance with the present invention;

Figure 2 shows flowcharts illustrating the general steps performed by the document-based database management system when used to develop a model of a business and corresponding data schema, create documents using the data schema and populate database tables using the created documents;

Figure 3 is a screen display of an entity page presented to a developer during creation of a business model using a modeller forming part of the document-based data management system of Figure 1;

Figure 4 is a screen display of the entity page of Figure 3 together with a new detail entity form;

Figure 5 is a screen display of the entity page together with an attribute window;

Figure 6 is a screen display of the entity page presenting a created model;

Figure 7 is a screen display of the entity page together with a model entity category form;

Figure 8a is a screen display of the entity page together with a data edit attribute form;

Figure 8b is a screen display of the entity page together with a calculation edit attribute form;

Figure 8c is a screen display of the entity page with a lookup edit attribute form;

Figure 9a is a screen display of the entity page together with a completed data edit attribute form;

Figure 9b is a screen display of the entity page together with a completed calculation edit attribute form;

Figure 9c is a screen display of the entity page together with a completed lookup edit attribute form;

Figures 10a and 10b are screen displays of the entity page together with attribute list and attribute count windows;

Figure 11 is a screen display of the entity page together with a registration code window;

Figure 12 is a screen display of the entity page together with a prefix entities window;

Figure 13a is a screen display of the entity page together with a unit manager window;

Figure 13b is a screen display of the unit manager window together with a group options window;

Figure 14 is a screen display of the entity page together with a choose database for model window;

Figure 15 is a flowchart showing the steps performed by the modeller during model and database table creation;

Figure 16 is a screen display of a document creation page presented to a developer during creation of documents using document tools forming part of the document-based database management system of Figure 1;

Figure 17 is an enlarged view of a document creation toolbar;

Figures 18a to 18d are screen displays showing the steps performed to place a data control in an active document presented in a document creation palette of the document creation window;

Figures 19a and 19b are screen displays of differently partitioned documents;

Figure 20 is a screen display of a document together with an entity link list window and entity link editor window;

Figures 21a to 21c are flowcharts showing the steps performed by the development tools during document creation;

Figure 22a shows a document-centered interface presenting a drilling startup document;

Figure 22b shows a document-centered interface presenting an AFE cost estimates document;

5 Figure 22c shows a document-centered interface presenting a rental sheet document;

Figure 22d shows a document-centered interface presenting a bit and pump details document;

Figure 23a shows a document-centered interface presenting a daily

10 drilling report document;

Figure 23b shows a document-centered interface presenting a daily drilling report document with a time/date range;

Figure 23c shows a document-centered interface presenting a daily cost sheet document;

15 Figure 23d shows a document-centered interface presenting a hole, casing and cementing details document;

Figure 23e shows a document-centered interface presenting a casing tally document;

Figure 23f shows a document-centered interface presenting a vendor

20 work order document;

Figure 24a shows a document-centered interface presenting a BHA equipment list document;

Figure 24b shows a document-centered interface presenting a formation and survey details document;

25 Figure 24c shows a document-centered interface presenting a cost control document;

Figure 25a shows a document-centered interface presenting a rig photo document;

Figure 25b shows a document-centered interface presenting a lease

30 condition photo document;

Figure 26 shows a document-centered interface presenting a completed job startup document;

Figure 27a shows a document-centered interface presenting a daily report document;

Figure 27b shows a document-centered interface presenting an equipment document;

5   Figure 27c shows a document-centered interface presenting a downhole schematic document;

Figure 27d shows a document-centered interface presenting the downhole schematic document with a metagram toolbar exposed;

Figure 27e shows a document-centered interface presenting the

10   downhole schematic document with a drill-down control exposed;

Figure 27f shows a document-centered interface presenting a wellhead schematic document with a drill-down control exposed;

Figure 28 shows a document-centered interface presenting a cost variance graph document;

15   Figure 29a shows a document-centered interface presenting an internal work order document;

Figure 29b shows a document-centered interface presenting the internal work order document with the halo and floating toolbar exposed;

Figure 29c shows a document-centered interface presenting a cost

20   variance report document;

Figure 30a shows a document-centered interface presenting a well status report document;

Figure 30b shows a document-centered interface presenting a depth vs. total days document;

25   Figure 30c shows a document-centered interface presenting a time breakdown graph document;

Figure 31a shows a document-centered interface presenting a drilling project cost variance document;

Figure 31b shows a document-centered interface presenting a project

30   costs graph document;

Figure 32a shows a document-centered interface presenting an Alberta Energy and Utilities Board (EUB) style WR-1 document;

Figure 32b shows a document-centered interface presenting an EUB style WR-2 document;

Figure 32c shows a document-centered interface presenting an EUB style WR-4 document;

5   Figure 33 is a schematic diagram of a distributed data network;

Figure 34 is a screen display of a communications interface; and

Figure 35 is a schematic block diagram showing telecommunications between a remote computer site and a primary control computer site.

10   **Detailed Description Of The Preferred Embodiment**

The present invention relates generally to data management and in particular to a document-based data management system and method for managing data over a distributed data network that includes one or more remote computer sites and a central primary computer site. Initially, a model of a business entity for which

15   data is to be managed is developed and is translated into a database table structure for a selected backend database into which data is to be stored. Documents are then created using attributes from a data schema automatically derived from the model. During document creation, data control types are selected and attributes in the data schema are dragged and dropped into an active document using a computer pointing

20   device. A data control of the selected type is created in the active document for each attribute dropped into the document. The data controls are automatically hooked up to the database through objects. Information entered into the data controls automatically populates the database tables through the objects. As a result, created documents are used to enter, edit and/or delete data in the database and view, print

25   and query data stored in the database.

Documents are grouped and are available throughout the distributed data network via workviews. The workviews include user-specified work folders containing user-specific documents. The documents are used to populate the database with data and thereby create records. Records are collected from remote sites of the

30   distributed data network and the records are reconciled with a central database at the primary computer site. Reports can be generated using documents that are hooked up to the databases in the distributed data network.

Turning now to Figure 1, a schematic of the functional components of the document-based database management system in accordance with the present invention is shown and is generally identified by reference numeral 100. As can be seen, the document-based database management system includes a set of development

5   tools 102 that are accessed by a modeller 104, a packager 106, workview tools 108 and document tools 110. The document-based database management system 100 also includes graphical user interfaces to allow developers and end-users to interact with the functional components of the system 100. In the preferred embodiment, the document-based data management system is written in Inprise's Delphi Enterprise

10   that is distributable across standard Windows '95, '98 and NT platforms. As is well known, Delphi is object-oriented software that facilitates Windows development using drag-and-drop visual programming. Delphi provides excellent database support, strong object-orientation and high productivity.

Modeller 104 allows a developer to create a model 111 of an entity

15   such as a business, organization or other entity of interest (hereinafter referred to as "business entity"). The modeller 104 also translates the model into a database table structure for a selected backend relational database 114 as will be described.

Packager 106 includes an export/import function and creates database packages 115 that are to be conveyed to another computer site 116 over a

20   telecommunications link 118. Database packages 115 received by the packager 106 120 are unpackaged and are then conveyed to the linking engine 112, which in turn reconciles the data in the database packages with the data in the database.

Workview tools 108 and document tools 110 allow the developer to create user-specific workviews that contain user-specific documents 126. The

25   document tools 110 create objects that are used by the linking engine 112 to live-link data controls in the documents 126 to the database 114. In this manner, data entered into the fields of the data controls automatically populates the database 114. Thus, data visible in a document through a data control is in the database 114. In this particular embodiment, the linking engine 112 interfaces with ODBC or SQL backend

30   databases such as Paradox, Oracle, Interbase, and SQL Server to name a few. Those of skill in the art will however appreciate that the linking engine 112 may be designed to interface with other databases if desired. The documents 126 also interact with

other applications 120 via an OLE interface 122. The general overall operation of the document-based database management system 100 will now be described with reference to Figure 2.

Initially, the developer uses the modeller 104 to develop a model that
5   represents the business entity and its information requirements. The model is developed visually using entity-relationship diagrams. Using the modeller 104, an entity within the business entity is created and displayed on a model creation palette and attributes are assigned to the entity. A second entity is then created or dragged and dropped onto the model creation palette and its relationship to the first entity is
10   defined graphically. This process is performed until a complete visual model 111 of the business has been created (blocks 140 and 142). Once this has been done, a table generator in the modeller 104 is conditioned to translate the visual model 111 into underlying hierarchical database tables (i.e. the data schema) that correspond to the model for the selected backend database 114 (block 144).
15   After the database tables are created, the developer uses the document tools 110 and the data schema to create documents 126 for the end-users in the business. During this stage, the data schema is presented in a window adjacent an active document. Data control types are selected and attributes from the data schema that are to form part of the active document are dragged and dropped into the active
20   document (block 154). Data controls of the selected types and associated with the attributes dropped into the active document are automatically created and are live-linked to the backend database through objects. This document creation process is continued until a full suite of documents is created that services all of the end-users of the business.
25   The development tools 102 include a full suite of data-aware control icons including for example: columns, data fields, graphs, sketches, pictures etc. so that various types of data controls can be included in each document. The workview tools 108 provide tools to allow the developer and end-users to control the overall manner in which documents are presented to end-users. As a result, the development
30   and workview tools allow the visual appearance of the documents and the manner and types of data controls in the documents to be effectively tailored to suit the needs of the end-users for whom the documents are created. Since the documents 126 are

created specifically for end-users, the documents reflect the end-users' views of the business model and are the users' gateways to the business model. The documents therefore control who is able to access what corporate data and how the corporate data is accessed.

5   Once the documents are created, the documents are arranged within workbooks according to user work requirements. The database management system 100 including the model, the database tables and the various workbooks are then installed on the computer sites throughout the distributed data network of the business entity. The documents are then used to populate the database with data. Since the

10   data controls in the documents are live-linked to the backend database 114 by the linking engine 112, as data is entered into the fields of the data controls (blocks 160 and 162), the linking engine 112 automatically populates the tables of the backend database 114 with the entered data (block 164). Records stored in the databases at remote computer sites that include new data are packaged by the packager 106 and are

15   conveyed over a telecommunications link 118 to the central primary computer site of the business entity. The packages are then imported by the packager 106 at the central primary computer site and are depackaged. Once depackaged, the records are reconciled with the central database by the linking engine 112. In this manner, all information entered into databases throughout the distributed data network can be

20   reconciled in the database at the central primary computer site. This information can then be effectively used by users to make educated business decisions.

        A specific embodiment of the document-based database management system 100 in accordance with the present invention will now be described more fully with reference to Figures 3 to 35. In this embodiment, system 100 manages data

25   within an Oil Company.

Model Creation

        Initially, the document-based database management system 100 is used by the developer to develop a visual model of the business entity and then translate

30   the model into a hierarchical database structure corresponding to the model for the selected backend database 114. The model creation and translation process is described below with particular reference to Figures 3 to 15.

Graphical User Interface

The graphical user interface allows the developer to interact with the modeller 104 and develop a visual model of the business entity starting from the top

5   down. When the developer accesses the modeller 104, the graphical user interface presents an entity page 170 to the developer as shown in Figure 3. A plurality of drop-down menus are provided near the top of the entity page and can be opened by the developer by selecting the appropriate drop-down menu using a computer pointing device. Below the drop-down menus is a toolbar 172 including a plurality of

10   selectable icons. The icons include a generate database tables icon 172a, an add detail entity icon 172b, an add derived entity icon 172c, a make entity a master of another icon 172d, a make entity a base of another icon 172e, a delete entity icon 172f and a disconnect all relationships for this entity icon 172g.

Below the toolbar 172 is a model creation palette 174 on which the

15   model is visually created using the icons or corresponding drop-down menu options. With a blank model creation palette 174 opened, the visual model of the business entity is created in a top-down manner by establishing the entities within the business entity (in this case the Oil Company), attributes assigned to the entities and the relationships between the entities.

20   To begin creation of a visual model, the add entity detail icon 172b is selected resulting in a pop-up blank entity detail form 176 being opened as shown in Figure 4. The blank entity detail form 176 includes an entity name field 176a, a category field 176b, a database name field 176c and a remote deletions field 176d. The entity name field 176a holds the name assigned to the entity. The category field

25   176b allows the entity to be grouped with other entities so that groups of entities can be moved within the model as a whole. The database name field 176c allows a database to be selected that is different from the database in which the rest of the model data is stored. The remote deletions field 176d limits or denies deletions received from remote computer sites as will be described.

30   When the entity detail form 176 is completed and saved, the entity is added to the model and is represented on the model creation palette 174 by a rectangular label bearing the name of the entity (see Figure 6). Once the entity is

-16-

added to the model, the entity is assigned a primary key attribute (Xprime), a last altered date (Xaltered) and a last arrived date (Xarrived). These attributes and dates are used to support a low-bandwidth distributed database as is described below.

The primary key attribute acts as a modified "auto-increment" and is

5   the sole identifier of a record of data throughout the entire distributed database. The primary key attribute is used to keep different records from overwriting one another. To inhibit different records from overwriting each other, a block of key numbers is pre-allocated to each different user and is termed the user license. In this particular embodiment, the user licence grants a user 2^8 keys with which to enter data into the

10   database. The primary key attribute is a 32-bit integer that includes an upper part and a lower part. The upper part of the integer is the user licence number. The lower part of the integer is an auto-increment, and increments by one for each new record created by the user for a given entity within the model, until the upper limit is reached. When the upper limit is reached and the user attempts to add a new record to

15   the database, the user is presented with a message stating "This user license has completely run out of keys". The user can then be assigned a new unique user license allowing the user to continue to create new records.

Since auto-increment field behavior is inconsistent amongst databases (for example, Paradox only lets fields begin from one, and does not let different users

20   working on the same database start from different spots), and since the auto-increment values are typically only available after a record has been posted to the central database, the linking engine 112 keeps track of the auto-increments on its own, and keeps the primary keys as simple long integers. In case of an improper shutdown, a value that is zeroed only when the database management system 100 exits can

25   indicate that a "recovery" needs to be done. During a recovery, the linking engine 112 finds the highest value in the pre-allocated key range that already exists, and auto-increments from there as new records are created.

The number of primary keys allocated for each user license is per entity and not all entities together. This is permissible since the records only have to

30   be unique across the distributed database in each table, not across all tables. As a result, the limiting factor for a given user license is the largest number of records entered for any entity.

The last altered date includes a date/time field and acts as a "stamp" to allow modifications to particular records to be tracked. When records are received by the central primary computer site from remote computer sites, the last altered date allows the linking engine 112 to determine if the data is newer than entries in the

5   central database, in which case the data is applied, or whether the data is older or the same date as entries in the central database, in which case the data is ignored.

The last arrived date represents the time that a given piece of data arrives at the central database. For records created or modified on the central primary computer site, the last altered date and the last arrived date assigned to the records are

10   the same. For records that are received by the central primary computer site from a remote computer site, the last arrived dates are set to the time at which the records are unpackaged by the packager 106. Thus, the differences between the last altered dates and last arrived dates indicate the length of time the records have been "out" of the central database. The last arrived dates are used when it is desired to search a

15   database for new records during a print new or export new task as will be described.

When an entity is added to the model, an attribute window 180 is also opened for the entity as shown in Figure 5. The attribute window 180 displays the names of attributes assigned to the entity and includes an Add button 180a, an Add Calc button 180b, an Add Lookup button 180c, and Delete and Close buttons as will

20   be described. The attribute window 180 also displays the primary key, the last altered date and the last arrived date.

Once the first entity within the business entity has been created and displayed on the model creation palette, additional related entities can be added to the model. This is achieved by highlighting the first entity and selecting either the add a

25   detail entity icon 172b or add a derived entity icon 172c. Once either the add a detail entity icon or add a derived entity icon is selected, a new blank entity detail form 176 for the entity is opened. When the fields in the entity detail form 176 are completed and saved, the entity is assigned a primary key attribute, a last altered date, a last arrived date and a master key to its master entity. The entity is also represented on the

30   model creation palette by a rectangular label bearing the display name of the entity. In this case, since the entity is related to another entity, the label is joined to the label of its master entity by a connection that visually indicates the relationship between the

entities. In the present embodiment, "one to many" relationships between entities are denoted by blue lines terminating in circles while "is derived from" relationships between entities are denoted by green lines beginning with triangles.

Figure 6 shows an entity page 174 visually presenting a completed
5   model 202 for the Oil Company. In this particular example, the model 202 includes a corporation entity 210, a division entity 212, an oilfield entity 214, an offshore_oilfield entity 216, a gen_cost entity 218 and a cost_code entity 220. The connections between the entities visually indicate whether an entity adds detail to an entity above it or is derived from an entity above it. As can be seen, in this example,
10   the division entity 212 and the cost_code entity 220 add detail to the corporation entity 210. The oilfield entity 214 adds detail to the division entity 212. The offshore_oilfield entity 216 on the other hand is derived from the oilfield entity 214 (i.e. the offshore_oilfield is an (IS-A) oil field).

Although Figure 6 shows a visual model for the Oil Company, those of
15   skill in the art will appreciate that the size and complexity of the model 202 has been reduced for ease of illustration. A complete visual model 202 for an Oil Company would typically include a number of corporate divisions, each responsible for a number of oil fields. The model 202 would also typically include entities associated with other physical categories within the Oil Company such as for example, the
20   seismology department, the accounting department etc. to name a few.

Changes between the entity-relationships in the visual model can be made to selected entities by highlighting the rectangular labels associated with the entities of interest and selecting either the make entity a master of another icon 172d, the make entity a base of another icon 172e, the delete entity icon 172f or the
25   disconnect all relationships icon 172g. If an entity is selected and the make entity a master of another icon 172d is selected, selecting a target entity will make the selected entity the master entity of the target entity thereby changing the relationship between the two entities. In this case, the visual model is updated to show the new master relationship between the entities. Similarly, if an entity is selected and the make
30   entity a base of another icon 172e is selected, selecting a target entity will make the target entity a derived entity of the selected entity thereby changing the relationship between the two entities. Again the visual model is updated to show the new

relationship between the entities. If an entity is selected and the delete icon 172f is selected, the selected entity and its relationships with all other entities are deleted from the model. If an entity is selected and the disconnect all relationships icon 172g is selected, all relationships between the selected entity and other entities in the model

5   are removed allowing new relationships with entities to be established using the make entity a master of another icon 172d and make entity a base of another icon 172e.

The visual appearance of the model can be enhanced by opening a model entity category window 230 via a category option in the view drop-down menu as shown in Figure 7. The category window 230 allows the visual appearance of the

10   rectangular labels associated with all of the entities in a selected category to be adjusted. As will be appreciated, entities are categorized through information entered into the category fields 176c of the entity detail forms 176. In this manner, the rectangular labels associated with entities of different categories can be colored differently making it easier to follow visually the model and differentiate between

15   model sections containing groups of entities. For example, if the model of the Oil Company includes entities within the accounting department, entities within the seismology department, and entities within each of the corporate divisions, the entities within each of these categories can be differentiated by color.

Anytime during the entity creation process, attributes can be assigned

20   to the entities. To assign an attribute to an entity, the attribute window 180 associated with the entity is used. Depending on the type of attribute to be assigned to the entity, the Add button 180a, Add Calc button 180b or Add Lookup button 180c is selected. Selection of one of the buttons opens a corresponding blank edit attribute form 240 including a number of data fields to be completed.

25   Selection of the Add button 180a opens a blank data edit attribute form as shown in Figure 8a. As can be seen, the data edit attribute form 240 includes an actual name field 240a, a kind field 240b, a unit group field 240c, a tag field 240d, a display name field 240e, a size field 240f, a decimal field 240g, an entity field 240h and a hint field 240i. The edit attribute form 240 also includes an Add another button

30   242, a Pick list button 244, a Cancel button and an Ok (save) button.

The actual name field 240a stores the name of the attribute as it appears in the database. The kind field 240b indicates the type of data such as for

-20-

example, text, integer, floating point etc. that is to be stored in the database in relation to the attribute. The available data types can be exposed in a menu by selecting the arrow trigger associated with the kind field 240b. The display name field 240e stores a text name for the attribute as it is to appear in attribute window 180 and as it appears
5   in a data schema window 340 (see Figure 16) that is opened during document creation. The size field 240f determines the number of characters or digits allocated to the attribute. The decimal places field 240g only becomes active when the kind field 240b stores a floating-point data type. When active, the decimal place field 240g stores a number that determines the number of decimals allocated to the attribute.
10   The hint field 240i allows notes to be inserted that relate to the attribute. The notes in the hint field 240 are displayed in a hint box on documents, which include data controls relating to the attribute, when a cursor is brought into proximity of the data controls. The tag field 240d can be selected to provide external links to files and to establish special treatment of attributes. The unit group field 240c
15   holds the unit type of the attribute such as depth, diameter etc. as set up in a unit manager as will be described. The unit group field 240c is available only for floating point data types.

            When the data edit attribute form 240 is completed and saved, the display name of the attribute is shown in the attribute window 180. Selection of the
20   add another button 242 saves the current data edit attribute form and opens a new blank data edit attribute form 240. Selection of the pick list button 244 allows a pick list of choices to be added to the attribute. When a document includes a data control corresponding to the attribute with a pick list, the field of the data control includes a trigger that allows the choices in the list to be displayed and an option from the list to
25   be selected for entry into the field of the data control.

            Figure 9a shows a completed data edit attribute form for a CORP_NAME attribute assigned to the corporation entity 210. As mentioned above, this form is opened in response to highlighting of the Corporation Name attribute in the entity window 180 and selecting the Add button 180a.
30   Selection of the Add Calc button 180b opens a blank calculation edit attribute form 250 as shown in Figure 8b. The calculation edit attribute form 250 is similar to data edit attribute form 240 but includes a script button 252 instead of a

pick list button. Selection of the script button 252 opens a script editor window 260 that allows a script to be written to carry out business rules on the data specified in the script. The script can be saved to an external file by selecting the external radio button 260a and entering a file name in the file name field 260b or can be saved

5   internally within the model by selecting the internal radio button 260c.

Figure 9b shows a completed calculation edit attribute form 250 and a completed script editor window 260 for a Difference attribute listed in the attribute window 180 associated with the division entity 212. In this particular example, the division entity 212 includes a Division Name attribute, a Division Budget attribute, a

10   Division Expenditures attribute, a Difference attribute and a master key to the Corporation attribute in addition to its primary key, last altered date and last arrived date fields.

Selection of the Add Lookup button 180c opens a blank lookup edit attribute form 266 as shown in Figure 8c. The lookup table edit attribute form 266 is

15   basically the same as the data edit attribute 240 form except the entity field 226h is active. Selecting the arrow trigger associated with the entity field 226h lists the other entities in the model. An entity can be selected from the list and entered into the entity field. In this manner, the attribute targets or "looks up" the attribute of another entity. The rectangular labels associated with entities that include lookup attributes

20   point to the entities associated with the targeted attributes by means of dotted purple arrows.

Figure 9c shows a completed lookup edit attribute form 266. In this particular example, the Amount attribute targets the Cost_Code entity and stores the primary key of the Amount attribute assigned to the Cost_Code entity. Figure 6

25   shows the purple dotted line 268 extending from the Gen_ Cost entity and pointing to the Cost_Code entity.

As will be appreciated, various attributes can be assigned to the entities in the model by opening edit attribute forms via the Add, Add Calc and Add Lookup buttons in the attribute windows 180.

30   During model creation or after the complete visual model has been created, information concerning the model can be visually displayed through selections in the drop-down menus. For example, all attributes assigned to entities in

-22-

the model can be listed in a pop-up window 270 by selecting an attribute option in the drop-down view menu as shown in Figure 10a. The header of the pop-up window 270 indicates the number of attributes in the model and a count of the data attributes and a count of the calculation attributes. In Figure 6, the model 202 includes twelve

5   (12) attributes, eleven (11) of which are data attributes and one of which is a calculation attribute. The type and location of the attributes in the model can also be presented in a text window 272 by selecting an attribute count option in the drop-down special menu as shown in Figure 10b. The ability to determine the number and types of attributes in the model 202 plays an important role in the licensing of the

10   document-based database management system 100 to licensees as will be described.

The graphical user interface also allows a registration code to be applied to the modeller 104 via a registration code window 274 as shown in Figure 11. The registration code window 274 is opened by selecting a registration option in the drop-down special menu. The registration code can be used to set a time limit on

15   the usage of the modeller for demonstration purposes and to limit access to the modeller 104.

The graphical user interface also allows prefixes to be assigned to models via a prefix entity window 276 as shown in Figure 12. The prefix entity window is opened by selecting a prefix option in the drop-down special menu.

20   Prefixes are used to keep the database table names for different models separate. This is important when dealing with databases such as Oracle, which has a system where tables for different databases are kept in the same database space.

The graphical user interface also includes a unit manager window 280 as shown in Figure 13a. The unit manager window 280 is opened by selecting a unit

25   manager option in the drop-down view menu. The unit manager window 280 includes three areas, namely a family area 282, a unit area 284 and a group area 286. The family area 282 displays a list of the unit families, the unit area 284 displays a list of units and the group area 286 displays a list of unit groups. New family and delete buttons 282a and 282b are associated with the family area, new unit and delete

30   buttons 284a and 284b are associated with the unit area and new group and delete buttons 286a and 286b are associated with the group area. These buttons allow entries to be added to and removed from the lists.

Each unit family in the list includes a family of units that measure the same parameter such as for example, length, area, pressure, volume etc. When a unit family in the family area 282 is highlighted, the available units associated with the highlighted family are listed in the unit area 284. For example, the length family

5   includes units such as meters, kilometers, miles, inches etc.

Each unit group in the list represents a semantic grouping of units. Selecting a group in the list opens a group option window 288 as shown in Figure 13b. As can be seen, group options window 288 includes a group name field 288a, a stored field 288b and a display field 288c. Group name field 288a identifies the

10   selected group. Stored and display fields 288b and 288c allow the stored and displayed units to be set to desired units. In this manner, units for a group can be displayed in one unit yet stored in the database in a different unit. In this case, the group option window 288 can be used to ensure that all units in the group are stored across the distributed data network in the same units while allowing for presentation

15   in different units. This is particularly advantageous when different areas of a business are distributed internationally and operate in different units (i.e. metric vs. imperial).

Once the complete visual model has been created, including all entities and attributes, the backend database type is selected by opening a database window 290 via a database option in the drop-down file menu as shown in Figure 14. The

20   database window 290 presents a list of database aliases from which the developer may select the backend database. Once the backend database has been selected, the database tables for the model that are compliant with the selected backend database can be generated by selecting the generate database table icon 172a. Selection of the generate database table icon 172a causes the table generator of the modeller 104 to

25   generate the appropriate hierarchical database tables for the selected backend database using the entities, attributes and entity-relationships within the model. As mentioned previously, the hierarchical database corresponds to the model so that records stored in the database tables are stored in a manner that tracks the structure of the business entity.

30   If it is desired to update an existing model, the model can be opened through the drop-down file menu and presented on the model creation palette 174. The model can be updated. Once updated, the generate database table icon 172a can

be selected to generate new database tables for the updated model that are compliant with the selected backend.

Object and Database Table Creation

5   When the generate database table icon 172a is selected, the modeller 104 selects an entity from the model (see block 200 in Figure 15) and creates a table object for the entity that hooks up to the selected backend database (block 202). The modeller 104 then creates field definitions from selected properties of the attributes associated with entity for the table object (block 204). Field definitions are then

10   added to the table object for the entity's master relationship (block 206). Field definitions for distributed database support using the primary key, last altered date and last arrived date fields are also added to the table object (block 208). Index definitions for master and primary keys are then added to the table object (block 210). At this stage, the modeller 104 asks the driver of the selected backend database to

15   create a database table for the entity using the table object (block 212). The above process is repeated for each entity in the model until database tables for all entities in the model have been created (block 214). As will be appreciated, since the database tables are generated using objects that correspond to entities within the model and which include field definitions corresponding to the attributes and relationships

20   between the entities, the database tables are hierarchical and follow the model. As the database tables are created, a data schema following the model is also created. The data schema is used to create documents as will be described.

          The table generator also creates two additional tables, namely an XENTITY table and an XDELETE table that are used to support remote deletions as

25   will be described. The naming convention used for these tables is chosen to avoid collisions with tables associated with entities in the model. The XENTITY table lists all of the entities in the model by number. The number of the entity is used as the primary key to avoid having the entity referred to by different records across the distributed data network. This reduces the complexity of processing deletions, since

30   all deletions for a given entity are kept under the same master record. The XDELETE table stores lists of the primary keys associated with records for each entity that have

been deleted and the master entities the records used to belong to. Once this has been done, the model creation process is complete.

Document Creation

5   Following the creation of the model 202 for the business and the generation of the hierarchical database tables, the development tools 110 and the workview tools 108 are used by the developer to create documents to be used by end-users within the business. The documents are created for specific end-users, teams, departments etc. (generically referred to as "users") and are defined by the model of

10   the business entity. The documents are used to populate the backend database 114. Further specifics of the document creation process will now be described with particular reference to Figures 16 to 21c.

Graphical User Interface

15   The graphical user interface allows the developer to interact with the document tools 110 and workview tools 108 to create documents. When the developer accesses the document tools 110, the graphical user interface presents a document creation window 300 including a document creation palette 302 and a workview 304 as shown in Figure 16. A plurality of icons 306 are presented above

20   the workview 304 and include from left to right, a New Document Folder icon, a New Document icon, a New Communications icon, a New Scheduled Event icon, a New Named Event icon, a Print New Response icon, a Print Response icon and an Import Response icon. Selection of these icons allows an individual workview 304 to be configured by the developer.

25   The workview 304 presents a workbook 310, folders 312 within the workbooks and documents 314 within the folders. In this manner, the developer can create documents and group related documents in folders 312. The New Document Folder icon and the New Document icon allow folders and documents to be added to the workview 304. The New Communications icon allows a communications

30   component to be added to the workview. The Event icons allow the developer to establish actions that are to be performed automatically when certain events occur.

A blank active document can be opened in the document creation palette 302 by selecting the New Document icon and unlocking the document creation facilities instead of choosing an existing document. During document creation, a band toolbar 320 can be exposed by selecting the toolbar option in the drop-down

5   view menu. The band toolbar 320 includes a halo icon 320a, a header icon 320b, a body icon 320c, a body strip icon 320d and a footer icon 320e. Selecting one or more of the icons allows the developer to partition the active document into separate bands as will be described.

A document creation toolbar 330 can also be exposed by selecting the

10 document creation toolbar option in the drop-down view menu. The document creation toolbar 330 includes a plurality of data-aware control icons and is shown in Figures 16 and 17. The data-aware control icons on document creation toolbar 330 in this particular example are arranged in groups and include data control entry icons, other icons, title/login icons and format icons.

15 The data control entry icons include a regular mouse mode icon 330a, a data entry icon 330b, a graph icon 330c, a column icon 330d, an X-box icon 330e, a light icon 330f, an X-box other icon 330g, a light other icon 330h. Selection of the mouse mode icon 330a, cancels other modes and places the cursor in a regular mouse mode. Thus, the cursor can be moved across the document but cannot be used to drag

20 new attributes from the data schema window 340 onto the document. The data entry icon 330b allows a data control having single data field to be dropped onto the document being created. The graph icon 330c allows an area in the document to be drawn out for a graph and allows attributes to be dropped into the drawn out area. The column icon 330d allows a column-type data control for an attribute to be

25 dropped onto the document. The X-box icon 330e allows a "check-box"-type data control for an attribute to be dropped onto the document. The light icon 330f is similar to the X-box icon and allows a light-type data control to be dropped onto the document. The X-box and light icons 330e and 330f can be used to represent logical (true or false) fields or to indicate choices from a list of options. In the case of the

30 check-box-type data control, the convention used is typically "checked" for true and "blank" for false. In the case of the light-type data control, the convention typically used is "bright" for true and "dim" for false. The X-box other icon 330g allows a

to tell the linking engine 112 how to treat records when attributes from a derived

entity are placed on the document from the base entity. The outer join icon 336d is

used to instruct the linking engine to display all of the records while the inner join

icon 336e is used to instruct the linking engine to display only those records with a

5   derived part in them.

         The data schema window 340 can be opened by selecting the schema

view option in the drop-down view menu. The data schema window 340 presents the

data schema of the model including the entities and attributes in the model. Symbols

342 accompany the entities and attributes to identify visually the type of attribute (i.e.

10   data, calculation or lookup) and the nature of the attribute (i.e. text, numeric, logical,

graphical etc.). A toolbar 344 is provided at the top of the data schema window 340

and includes a plurality of selectable icons that allow the manner in which the data

schema of the model is presented to be controlled. Specifically, the toolbar 344

includes a tree icon 344a, a sort attribute by name icon 344b, a sort attribute by type

15   icon 344c, a sort attribute by entity icon 344d, a show/hide data attributes icon 344e

and a show hide calculated attributes icon 344f.

         When the tree icon 344a is selected, the data schema of the model is

presented in a tree form as shown in Figure 16. When the sort attribute by name icon

344b is selected, the attributes in the data schema are presented in alphabetical order

20   with their associated entities in parentheses. When the sort attribute by type icon

344c is selected, the attributes in the data schema are grouped according to type.

When the sort attribute by entity icon 344d is selected, the attributes in the data

schema are.grouped according to entity. When the show/hide data attributes icon

344e is selected, data attributes in the data schema are hidden from view. When the

25   show/hide calculated attributes icon 344f is selected, calculation attributes in the data

schema are hidden from view.

         During document creation, the band toolbar 320 is used to partition the

active document into the desired bands. The document creation toolbar 330 is then

used to select the type of data control to be inserted into the document. Following

30   this, the attribute of interest is dragged from the data schema window 340 and

dropped into the appropriate band of the active document being created. Once an

attribute from the data schema window is dropped into the active document, a data

control of the type determined by the selected data control icon is created on the document and is live-linked to the appropriate database table. The size of the data control field is determined by the information entered into the kind, size and decimal place fields 240b, 240f and 240g of the edit attribute form when the attribute was

5   created. For example, turning to Figures 18a to 18d, the placing of a data control onto an active document is shown. In this case, the data entry icon 330b is selected and the wave height attribute in the data schema window 340 is highlighted (see Figure 18a). Using the computer pointing device, the wave height attribute is dragged from the data schema window to the desired location on the active document (see Figures 18b

10   and 18c). The wave height attribute is then dropped into the active document resulting in an active wave height data control 350 being created automatically in the document.

This attribute drag-and-drop process is continued until all of the bands of the active document include the desired number and types of data controls of

15   interest. When a data control for an attribute is created on the active document, it is necessary for that document also to include at least one data control for an attribute of each entity above it so that a link to the top of the data schema hierarchy exists within the document. For example, in Figure 16, if the wind attribute of the offshore_oilfield entity is dragged from the data schema window and dropped into the active document,

20   in order to complete the document, an attribute from the division entity and an attribute from the corporation entity must also be dragged from the data schema window 340 and dropped into the active document. Although each document must include data controls for attributes, which follow the entity chain to the top of the data schema hierarchy, often it is not desired to display all of these data controls in each

25   document.

The band toolbar 320 allows the active document to be partitioned to establish the data controls which are visible to users and to establish the manner in which documents are printed. As mentioned above, often it is not desired to display all of the data controls in a document. The halo band is designed to contain data

30   controls that are required for the document but which are not to be displayed or printed. The header and footer bands interact with the body strip band during printing of documents. During a print cycle, a new page is printed when the body strip band is

too large for the remaining printable area on a page or when the field of a data control in the header band or footer band switches. Printing of the body strip band repeats in the space between the header and footer bands.

Figures 19a to 19b show examples of documents that have been
5   partitioned differently using the band toolbar icons. In Figure 19a, the document is partitioned to include a halo band, a header band, a body band and a footer band. Data controls are included in each of the bands. In Figure 19b, the document is partitioned to include a halo band, a header band, a body strip band and a footer band.

During document creation, the developer may wish to view the links in
10 the document to determine how the data controls in the document hook up to the database. Figure 20 shows a document with an entity link list window 360 and an entity link editor window 362 opened. The entity link list window 360 is opened by selecting the entity links option in the drop-down view menu and the entity link editor window 362 is opened by selecting the edit button 360a in the entity link list window.
15 As can be seen, the entity link list window presents a list of the entity links while the entity link editor window provides specifics with respect to a selected entity link from the list displayed within the entity link list window. The entity link list window 360 and entity link editor window 362 function as a diagnostic tool to allow the developer to see how links are hooked up to the database and how the links behave.

20 Also, during document creation, the developer can assign user access requirements to the document that is displayed in the document creation palette by selecting an administrative document requirements option in the drop-down special menu and setting the desired user level. If this is done, only users with all of the requirements assigned to the document by the developer are permitted access to the
25 document.

Object Creation and Live-linking

As mentioned above, as soon as data controls are created on the active document, the data controls are live-linked to the database. In order to live-link the
30 data controls to the database, when an attribute is dragged from the data schema window and dropped onto the active document, the document tools 110 determine the type of data control selected for the dropped attribute based on the selected data-

-31-

aware control icon of the document creation toolbar 330 and the type of attribute and then create a control object for the selected attribute (block 408). The document tools 110 then ask the entity link owner for an entity link object that is hooked up to the entity to which the dropped attribute has been assigned (blocks 410 and 411). If an
5 appropriate entity link object does not exist, the document tools create an entity link object for the attribute (block 412).

Following this, an entity reference object is created (block 413) and the attribute is added to the entity link object. The process of adding the attribute to the entity link object creates an attribute reference object from the entity link object to the
10 attribute (block 414). A data source object from the entity link object is then hooked up to the created control object, which triggers the entity link object to generate appropriate SQL code (block 416). If the entity link object does not have a query contained in it, a query is created (block 418).

The entity link object then generates appropriate field objects for the
15 query (block 420). In the present embodiment, it is important that field object generation is performed by the entity link object and not the query since the business model may support data types that are not supported by the actual selected database.

The entity link object then polls the entity link owner object to determine if its associated entity has a master entity in the model (block 422). If the
20 entity has a master entity, the entity link object generates a master link to the entity link object corresponding to that master entity (block 424). When the entity link object of the master entity detects that an entity link object has linked to it, the entity link object of the master entity in turn polls the entity link owner object to determine if its associated entity has a master entity in the model. If so, the above steps are
25 performed and this process is repeated until an entity link object links to the entity link object of a master entity that itself has no master entity (block 426). As a result, master links are generated between entity link objects that correspond to the connections between their associated entities in the model.

Once master links between entity link objects are created, SQL code is
30 generated by the entity link object according to its fields and master links (block 428). The query is then opened (block 430). Opening of the query triggers the document tools to open the data source object, which in turn hooks the field of the data control

in the document up to the database (block 432). With the data control hooked up to the database by the data source object, the data control is live-linked to the database. As a result, data entered into the field of a data control is automatically entered into the database.

5   If a grid type data control (i.e. a column or graph) has been dropped onto the active document and another grid type data control is dropped into the same area of the active document, the adjacent grid type data controls are automatically linked. During this process, the grid control is automatically created at block 408. At block 410, if the appropriate entity link object does not exist, the document tools

10   inhibit the selected attribute from being dropped onto the active document. If the entity link object exists at block 411, blocks 414 to 432 are performed thereby to create the linked data control. The linking engine 112 manages the entity link owner and entity link objects.

15   **Workbook**

Once all of the document types for the business have been created, workbooks including user-specific documents are created for the business entity. Each workbook and associated documents are created for either a specific user, team, department etc.

20   Figures 22a to 32c show the documents in a workbook specific to oil well drilling 500. As can be seen, the workbook 500 includes a document selection panel 502 and an active document panel 504. The document selection panel 502 visually presents the folders in the workbook and the user-specific documents within the folders. Highlighting a document within a folder opens the highlighted document

25   in the active document panel 504. Once the document is opened, the user can interact with the document via a document-centred user interface as will now be described.

A floating toolbar 506 is shown in Figure 22b, although the toolbar may be docked above the document selection panel 502. The toolbar includes a number of icons divided into two groups 508 and 510 respectively. The icons of the

30   first group 508 define shortcuts and remain the same regardless of the document opened in the active document panel 504. The icons of the second group 510 also

define shortcuts but are document specific and change depending on the document opened in the active document panel 504.

In the example shown in Figure 22b, the toolbar 506 includes a First icon 508a, a Previous icon 508b, a Next icon 508c and a Last icon 508d in group 508
5　and a Calcs icon 510a and a Start New Day icon 510b in group 510. The icons in group 508 remain inactive until a user clicks into a data control in the opened document. When this occurs, the icons become active. Selecting these icons allows the user to step through records created using the documents.

The Calcs icon 510a is selected by default and when selected, enables
10　scripts associated with data controls in the document presented in the active document panel 504. If it is not necessary to view the results of calculations, the scripts can be inhibited by deselecting the Calcs icon 510a or by pressing the "F2" key. In this manner calculations are only performed on data in the database when the document is presented in the active document panel 504 and the associated scripts are enabled.

15　Since the database tables correspond to the business model, data is stored hierarchically in the database. If the model includes multiple branches stemming from the same entity, a document having a data control associated with an attribute of that entity can include a drill-down control associated with the attribute of that entity. Selection of the drill-down control using the "F3" key allows one of the
20　various branches to be selected. Once a branch is selected, data exposed in documents in the active document panel corresponds to data related to the selected branch. Thus, the drill-down control allows the hierarchical data to be exposed.

The workbook 500 in this example includes a variety of documents arranged in folders and sub-folders for an Oilfield of the Oil Company. As can be
25　seen, the document selection panel includes a workbook having a drilling folder, a completions folder, an accounting folder, a management folder, a guide folder and a communications folder.

The drilling folder includes four sub-folders namely, a job setup folder, a daily folder, an optional folder and a summary folder. The job setup folder includes
30　a drilling startup document, an AFE cost estimates document, a service company list document, a rental sheet document and a bit and pump details document.

-34-

The daily folder includes a daily report document, a daily cost sheet document, a hole, casing and cement document, a casing tally document, a vendor work order document and a well summary document.

The optional folder includes a BHA setup document, a BHA

5   equipment list document, a mud inventory document, a cement estimates document, a DST report document, a formation and survey document, a MACP document, a cost control sheet document and a plugback and abandonment document.

The summary folder includes a bit summary document, a rig photo document and a lease condition photo document.

10   The completions/workovers folder includes three sub-folders namely, a setup folder, a daily folder and a summary folder. The setup folder includes a job startup document and an AFE cost estimates document. The daily folder is closed, hiding the documents therein from view. The summary folder includes a variance report document and a variance report graph document.

15   The accounting folder includes an accounting work order document, an accounting cost control document, a cost variance document and a cost variance comments document.

The management folder includes well, project and company sub-folders. The well and company sub-folders are closed hiding the documents therein

20   from view. The project sub-folder includes a drilling project costs document, a drilling projects costs graph document and a drilled well summary by project document.

The guide folder includes a well status report document, a well drilling data document, a well incident data document, a well abandonment notification

25   document and a working interest owners document. As will be appreciated, these folders and user-specified documents allow information concerning an oil well from start to finish to be captured, edited, viewed, queried and printed.

Figure 22a shows the drilling startup document. As can be seen, the drilling startup document includes a number of data type data controls and associated

30   data fields relating to an oil well, its location and its elevation.

Figure 22b shows an AFE cost estimates document. This document includes a number of data type data controls, a calculated data control as indicated by

the shading, and linked column data controls. In this document, the toolbar 506 is in the floating condition and the Calcs icon 510a is selected enabling the script associated with the calculated data control.

Figure 22c shows a rental sheet document. This document also shows a number of linked column data controls. Some of the linked data controls are calculated data controls as indicated by the shading. The scripts associated with the calculated data controls are enabled. An information window is also shown.

Figure 22d shows a bit and pump details document. This document includes a number of linked column data controls and associated data fields relating to bit and pump details.

Figure 23a shows a daily drilling report document. This document includes a number of data controls and associated data fields related to the drilling status of a well for a particular day. A drill-down control is exposed to allow the particular day to be selected. Selecting a different day from the drill-down control exposes the drilling report data entered for that particular day. This document also includes calculated data controls as indicated by the shading. Since the fields associated with the calculated data controls display numbers, the associated scripts are enabled.

Figure 23b shows the daily drilling report document of Figure 23a with a time/date range. The time/date range is used as a filter to filter the daily drilling report data and is presented in a window overlying the daily drilling report.

Figure 23c shows a daily cost sheet document. This document includes data controls and associated data fields relating to daily costs incurred during drilling of the oil well.

Figure 23d shows a hole, casing and cementing details document. This document includes data controls and associated data fields relating to the oil well's hole, casing and cementing details. As can be seen in the data fields associated with the casing information, the calculated data control as indicated by the shading is disabled.

Figure 23e shows a casing tally document. This document includes data controls and associated data fields relating hole and casing details and adjustments to them and a total tally for the casing.

Figure 23f shows a vendor work order document. This document includes data controls and associated data fields relating to oil well work orders.

Figure 24a shows a BHA equipment list document. This document includes data controls and associated data fields relating to a BHA equipment list.

5   Figure 24b shows a formation and survey details document. This document includes data controls and associated data fields relating to oil well formation and survey details.

Figure 24c shows a cost control document. This document includes data controls and associated data fields relating to oil well cost controls and includes

10   linked column calculated data controls as indicated by the shading.

Figure 25a shows a rig photo document. This document includes data controls and associated data fields relating to an oil rig. Of interest, a bitmap image of the rig is presented in the rig photo data control.

Figure 25b shows a lease condition photo document. This document

15   includes data controls and associated data fields relating to the condition of the well property at the time it was leased. Of note, the lease photo data control presents a bitmap image of the land.

Figure 26 shows a completed job startup document. This document includes data controls and associated data fields associated with the completion of the

20   oil well.

Figure 27a shows a daily report document. This document includes data controls and associated data fields relating to daily information collected during the completion and workover of the oil well.

Figure 27b shows an equipment document. This document includes

25   data controls and associated data fields relating to equipment used during the well completion and workover. Of note, an equipment photo data control presents a bitmap image of equipment of interest.

Figure 27c shows a downhole schematic document. This document includes data controls and associated data fields relating to the downhole schematic of

30   the oil well.

Figure 27d shows the downhole schematic document of Figure 27c with a metagram toolbar exposed. The metagram tool bar allows the illustrated schematic to be edited as desired.

Figure 27e shows the downhole schematic document of Figure 27c with a drill-down control exposed. The drill-down control allows various stages of the downhole schematic to be presented on the document by selecting a record from the drill-down control.

Figure 27f shows a wellhead schematic document. This document includes data controls and associated data fields relating to the oil well wellhead schematic. A drill-down control is exposed. Selecting records from the drill-down control allow different wellhead schematics to be presented in the document.

Figure 28 shows a cost variance graph document. This document shows a bar graph illustrating the estimated and actual costs associated with the various stages of the oil well.

Figure 29a shows an internal work order document. This document includes data controls and associated data fields associated with an internal work order. Of note, this document includes an X-box data control.

Figure 29b shows the internal work order document of Figure 29a with the halo and floating toolbar exposed.

Figure 29c shows a cost variance report document. This document includes data controls and associated data fields relating to the cost variance during drilling. As can be seen, this document includes calculated data controls as indicated by the shading.

Figure 30a shows a well status report document. This document includes data controls and associated data fields relating to the status of the oil well. As can be seen, this document includes light-type data controls.

Figure 30b shows a depth vs. total days document. This document includes data controls and associated data fields relating to the depth of the well versus time in days. As can be seen, the well depth information is plotted on a graph as a function of time in days.

Figure 30c shows a time breakdown graph document. This document includes data controls and associated data fields relating to the oil well. As can be

seen, this document presents a pie chart showing the time breakdown of drilling related activities.

Figure 31a shows a drilling project cost variance document. This document includes data controls and associated data fields relating to cost variances in
5   the drilling project.

Figure 31b shows a project costs graph document. This document includes data controls and associated data fields relating to oil well project costs. The cost information is presented in a bar graph and is broken down into various categories.

10   Figure 32a shows an Alberta Energy and Utilities Board (EUB) style WR-1 document.

Figure 32b shows an EUB style WR-2 document.

Figure 32c shows an EUB style WR-4 document.

As will be appreciated, most of the documents in the workbook are
15   used to allow users to enter, edit and/or delete records in the database. All of the documents are used to present records and data. When a user is required to enter data in the database, in this example, to complete the daily drilling report document shown in Figure 23a, the user opens a blank daily drilling report document. The appropriate data can then be entered into the fields of the data controls on the document.

20   As data is entered into a field of a data control of the document, the data source object associated with the data control that is hooked to the database automatically populates the database with the data. As a result, entered data that is displayed in a field of the data control immediately becomes part of the database.

During the data entry process, a record for the data is created under the
25   appropriate entity and the record is assigned a primary key. The record is also assigned a last altered date, which identifies the user creating the record and the date and time the record was created. The record is also assigned a last arrived date. This process continues until the report document has been completed.

If a data control corresponding to the same attribute is within another
30   document, when that other document is placed in the active document panel, the field of the data control is automatically filled with the data in the database, since the data source object is hooked to the database.

Since the fields in the data controls are hooked to the database whenever a document is placed in the active document panel, if a document is in the active document panel and fields in the data controls therein are not filled with data, no data exists in the database for those fields.

If an existing document is opened, the fields in the data controls are automatically filled with the corresponding data in the database. If the user edits data in the fields, the corresponding records in the database are updated and the last altered dates and last arrived dates are updated accordingly.

Document Printing, Queries and Grafts

The document-centered interface includes a print option, which when selected, prints information in the hierarchical database from the top down. The document-centered interface also includes print cycle, print new and print EUB style options in the drop-down file menu. Selection of the print cycle option presents an under current item option and an on current item option.

If the under current item is selected, the document-centered interface examines the document in the active document panel and the highlighted data control in the document and cycles through and prints the data in the database beneath the level of the highlighted data control.

If the on current item is selected, the document-centered interface examines the document in the active document panel and the highlighted data control in the document and cycles through and prints the data set at the selected level of the highlighted data control.

If the print new option is selected, the linking engine filters the database from the top of the model through a skeleton. The skeleton is a bare representation of the selected data that is applicable across the model, in this case, the new data since the last print new. During the filtering process, a skeleton following the model is built from data in the database that is new since the last print new selection.

If the print EUB style option is selected, the documents are printed in a form that satisfies Alberta Energy and Utilities Board requirements.

The document-centered interface includes a move option in the drop-down edit menu. Selection of the move option presents a start option. Highlighting a data control in the document in the active document panel and selecting the start option allows all of the data in the database in the branch associated with the

5   highlighted data control, to be moved to another branch. To complete the move, the user navigates the data control to the master entity to which the data is to be grafted and selects a finish option.

The document-based database management system also includes a query tool to search the database for desired information to be searched. The query

10   tool performs flat searches of the database that follow the entities of the business model for each parameter of the search. The results of the flat searches are then built up into a skeleton that is used to filter the data to yield only the pertinent search results.

15   **Distributed Environment**

As mentioned above, the workbooks are created for various users within the business entity. The documents within the workbooks are therefore particular to the users and include data controls that accept and present information of relevance to the users.

20   Once the workbooks are created for the business entity, the document-based database management system 100 including the model, the database tables and the workbooks (excluding the modeller and document creation tools) are installed on the computer systems throughout the distributed data network of the business entity. Users are able to access the workbooks of relevance to them and use the documents

25   therein to interact with the databases.

Figure 33 shows a typical distributed computer system environment in which the document-based database management system 100 is installed. As can be seen, the distributed computer system 520 includes a plurality of remote computer sites 522 and 524, two of which are shown for illustrative purposes, and a central

30   primary computer site 526.

Remote site 522 includes a computer 528 having a single workbook 500 loaded thereon. In this example, the workbook 500 includes two folders 500a.

One of the folders includes four documents 500b while the other folder includes three documents 500b. The documents in the workbooks are hooked to a database 530 including database tables corresponding to the model. The database 530 is a subset of the central database at the primary computer site 526.

5   The computer 528 communicates with a server 532 at the primary computer site 526 over a telecommunications link via its communications. The communications follow a multi-channel communications protocol as will be described. Specifically, the computer 528 communicates with a ground satellite dish 534 via a modem 536. Satellite dish 534 communicates with an orbiting satellite 538

10   that in turn communicates with a ground telephone communications system 540. The server 532 communicates with the ground telephone communications system 540 via its internal modem pool 542.

     Remote site 524 also includes a computer 528 having a single workbook 500 loaded thereon. In this example, the workbook 500 includes two

15   folders 500a. One of the folders includes two documents while the other folder includes one document. The documents in the workbook are hooked to a database 530 including database tables corresponding to the model. The database 530 is a subset of the control database at the primary computer site 526.

     The computer 528 at remote site 524 also communicates with the

20   server 532 at the primary computer site 526 over a telecommunications link via its communications. The communications also follow the multi-channel communications protocol. Specifically, the computer communicates with an radio frequency (RF) transmitter 546 via a modem 548. Transmitter 546 communicates with the ground telephone communications system 540, which in turn communicates

25   with the server 532 in the central primary computer site 526 via its internal modem pool 542.

     The central primary computer site 526 in addition to server 532, includes a print server 550 and a database server 552. A plurality of computers 554 are coupled to the servers via a network and similarly have workbooks 500 loaded

30   thereon. The database server 552 hosts the central database and the database tables generated by the modeller 104 during creation of the business model and stores the

data for the business entity entered through data controls in the documents of the workbooks across the distributed computer system.

As mentioned above, the remote computer sites 522 and 524 and the primary computer site 526 each include communications to establish a

5   telecommunications link and allow data to be exchanged over the telecommunications link. As can be seen in Figure 35, the communications is shown and includes a packetizer 600 to packetize outgoing data packages and to depacketize incoming data packets and a packet handler 602 to handle packets sent and received over the telecommunications link. The packet handler 602 follows a communications protocol

10   to reduce the amount of duplicate data that is transmitted during a telecommunications session as will be described. Although not shown, the communications also include the appropriate hardware and software to establish a telecommunications link with another computer site.

15   **Communications Interface**

When it is necessary to transmit records in a local database at a remote computer site to the central primary computer site 526 for storage in the central database, the export new option is selected from the drop-down file menu in the document-centered interface. Selecting the export new option causes the packager

20   106 to examine the last arrived dates associated with the records in the database to determine records that include data that is new since the last export to a particular computer site. The packager 106 in turn creates a file that includes the records (including the primary keys, last altered dates, last arrived dates and master keys). The file is then packaged and queued for future delivery.

25

Graphical User Interface

Selection of the communications module in the document selection panel opens a communications interface 700 in the active document panel as shown in Figure 34. The communications interface 700 allows packaged files at the remote

30   computer sites to be sent to the central primary computer site and allows updates from the central primary computer system to be sent to the remote computer sites. The communications interface includes a toolbar 702 at the top of the interface having a

plurality of selectable buttons including a Connect button 704, an Answer button 706, an Export button 708, a New button 710, an Edit button 712, a Delete button 714, a Queue Other Files button 716 and a Clear Queued Files button 718. Select Communications Device and Inactivity Timeout buttons 720 and 722 and Answer At

5   Program Startup, Auto-Send Waiting Files and Auto-Delete Sent Files check boxes 724 to 28 are positioned below the row of buttons. A list box 730, which is positioned below the buttons and checkboxes, lists the remote users to which the document-centered interface can connect.

Highlighting a destination computer site in the list box and selecting
10   the Connect button 704 causes the communications of the remote computer 528 to dial automatically the number of the destination computer site to establish a telecommunications link with the destination computer site.

Selecting the Answer button 706 places the communications of the remote computer into a call-waiting mode. If an incoming call is received when the
15   communications is in this mode, the incoming call is automatically answered.

The Export button 708 can also be used to cause the packager 106 to package all new records that have been created since the last time an export was performed by the user to a particular remote computer site, and queue the package for future sending.

20   A communications panel 740 is also included in the communications interface 700 and includes an Answer icon 742, a Connect icon 744, a Disconnect icon 746, a File Transfer Window Open icon 748, a Chat Window Open icon 750, a Connected Status indictor 752 and a Ringing Status indictor 754.

When a telecommunications link is established with a destination
25   computer site, the Connected Status indicator 752 is illuminated. When the Answer icon is selected to set the communications to answer incoming calls, the Ringing Status indicator 754 illuminates when an incoming call is received.

Selection of the File Transfer Window Open icon 748 opens a file transfer window 760. The file transfer window includes a list box 762 listing the
30   names of the packaged files that have been queued and the sizes of the files. The current file percentage transmitted and total percentage transmitted are shown by active progress bars 764 and 766 located below the list box. Select, Send and Abort

buttons 768 to 772 are positioned beside the list box. The number of bytes in the file that have been transferred out of the total bytes in the file, the elapsed time of the telecommunications link and the estimated time to complete the file transfer are displayed below the list box. The Select button 768 allows all of the packaged files in

5   the list box to be highlighted. Selection of the Send button 770 causes the packetizer 600 to packetize the highlighted packaged file or files one at a time and send the data packets to the packet handler 602 for transmission over the telecommunications link. The current file sliding bar indicates the percentage of the current file that has been transferred while the total sliding bar indicates the percentage of all data that has been

10   transferred to the remote location.

          When data is being received, a file transfer window 780 is opened and includes an active progress bar 752 to indicate the percentage of the file that has been received. Below the sliding bar, the file transfer window indicates the elapsed time of the telecommunications link, the estimated time to complete the file transfer, the

15   estimated total time of the telecommunications link, the name of the file being received and the number of bytes in the file that have been received out of the total bytes in the file.

          Selection of the Chat Window icon 750 opens a chat window 784 divided into incoming text and outgoing text boxes 786 and 788. Text entered into

20   the outgoing box is transmitted over the telecommunications link to the remote user. Responses from the remote user entered in the outgoing box of their chat window are transmitted over the telecommunications link and appear in the incoming text box of the chat window 784. The multi-channel communications protocol allows filed to be transferred between computer sites and chat conversations between users at the

25   computer sites to be carried out over the telecommunications link simultaneously.

Communications Protocol

          Prior to transmission, the packaged file is compressed and scrambled. Once this is done, the package is broken into packets by the packetizer 600 and the

30   packets are conveyed to the packet handler 602. The packet handler 602 in turn transmits the packets to the destination computer site. As packets are being sent, an entry is made in a table signifying that the packet is en-route to its destination and a

current transmission point within the file is maintained. When receipt of the packet is acknowledged by the destination computer site, the entry in the table is removed. If the communications link fails, the packet handler 602 determines the point in the file where transmission stopped together with the entries in the table representing

5   unacknowledged packets. When the communications link is re-established, the packet handler checks to determine whether the same file is to be transmitted. If so, the packet handler assembles the packets corresponding to the table entries and transmits the packets. The packet handler then continues transmitting the file packets from the point at which the transmission was stopped. In this manner, the packet handler 602

10   does not require transmission of data packets to be duplicated.

As the destination computer site receives the packets, the packets are held by the packet handler until the complete file is received. Once a complete file is received the packets are reassembled by the packetizer and the resulting package is decompressed and unscrambled.

15   At this time, the packager 106, for each record in the file, compares the primary key assigned to the record with the primary keys in the database. If an existing record exists in the database with the same primary key, the last altered dates of the two records are compared. If the transmitted record is newer, the record is stored in the database. In this case, the last arrived date is stamped with the current

20   time. Otherwise the record is ignored.

If no existing records exists in the database, the record is added to the database. In this case, the last arrived date is stamped with the current time.

If the transmitted file includes one or more record deletions, the packager 106 checks to see if the user who made the deletion was assigned deletion

25   rights via an entry in the remote deletions field 176c during model creation. If the user does not have deletion rights, the deletion is ignored.

If the user is assigned deletion rights, the packager 106 processes the deletion record-by-record. If there is no record in the database corresponding to the record to be deleted, the deletion is ignored. If a record is present in the database that

30   corresponds to the record to be deleted, the last altered data of the record in the database is compared with the last altered data of the deletion. If the last altered date

of the deletion is newer, the deletion is applied to the database and the XENTITY and XDELETE tables are updated.

The document-based database management system 100 allows documents to be tailored to meet the needs of the business of interest. Since the documents are live-linked to the database, the documents function as data input forms to the database and as data reports. The data controls in the documents determine the information that populates the database. Depending on the information needs of the business, the number of documents and the number and type of data controls in the documents may vary. Since modeller 104 and the document tools 110 allow the developer to determine the number and type of data controls and attributes in the model, the document-based database management system can be licensed to third parties and a royalty charged based on the number and type of attributes in the model and the number of attributes used across all of the documents.

Although a preferred embodiment of the present invention has been described with reference to the accompanying drawings, those of skill in the art will appreciate that variations and modifications may be made without departing from the spirit and scope thereof as defined by the appended claims.

What is claimed is:

1.   A document-based data management method comprising the steps of:
          developing a model representing a business entity's schema for which
5   data is to be managed; and
          creating documents using the model and live-linking the documents to
a database structure corresponding to said model so that data entered into the
documents populates the database structure.

10   2.   The method of claim 1 further comprising the step of translating the
model into said database structure for a database of interest.

3.   The method of claim 1 wherein said developing step includes the step
of using entity-relationship diagrams to develop said model.

15

4.   The method of claim 3 wherein said translating step includes the step
of generating database tables corresponding to said entity-relationship diagrams.

5.   The method of claim 1 wherein said documents are live-linked to said
20   database structure through document objects.

6.   The method of claim 5 wherein said document objects automatically
populate the database structure with data entered into said documents.

25   7.   The method of claim 6 wherein said documents include data controls
corresponding to attributes of said model and wherein said data is entered into fields
of said data controls.

8.   A workbook comprising:
30   a user interface including a document selection panel listing at least
one selectable document; and a document panel presenting a document selected from

said document selection panel, said document including at least one data control having a field to receive data; and

a linking engine live-linking said at least one data control to an underlying database, said linking engine automatically populating said database with
5   data entered into the field of said at least one data control.

9.   A workbook as defined in claim 8 wherein said linking engine is object-based and live-links said at least one data control to said database.

10   10.   A workbook as defined in claim 9 wherein said document selection panel includes a plurality of selectable documents.

11.   A workbook as defined in claim 10 wherein said selectable documents are held in folders listed in said document selection panel.

15

12.   A workbook as defined in claim 11 wherein each folder holds related user-specified documents.

13.   A workbook as defined in claim 10 wherein said user interface further
20   comprises a toolbar having a plurality of selectable icons, the functions of at least some said icons being user-definable.

14.   A workbook as defined in claim 13 wherein the functions of said at least some of the icons remain the same regardless of the document presented in said
25   document panel.

15.   A workbook as defined in claim 14 wherein the functions of other icons of said toolbar are also user-selectable, the functions of said other icons changing depending on the document presented in said document panel.

30

16.   A workbook as defined in claim 9 wherein each document includes a plurality of data controls, said plurality of data controls being live-linked to said database.

5   17.   A workbook as defined in claim 16 wherein at least one of said data controls includes a field presenting the results of a calculation.

18.   A workbook as defined in claim 17 wherein said calculation is carried out by a script.

10

19.   A workbook as defined in claim 18 wherein said script is saved within a model used to develop data tables within said database.

20.   A workbook as defined in claim 19 wherein said script can be toggled

15   between enabled and disabled conditions.

21.   A workbook as defined in claim 20 wherein the at least one data control that includes a field presenting the results of a calculation has a visual appearance different from other data controls.

20

22.   A workbook as defined in claim 9 wherein said database stores data hierarchically, said user interface further including a drill-down control to allow information associated with different branches in said database at a selected level to be displayed within a document presented in said document panel.

25

23.   A workbook as defined in claim 9 further including a selectable communications option, selection of said communications option opening a communications interface including controls selectable to initiate a telecommunications link with a remote site.

30

24.   A workbook as defined in claim 23 wherein said communications option is selectable from said document selection panel.

25. A modelling tool for creating database tables compliant with a backend database comprising:

    a user interface including a model creation palette and user-selectable controls to enable a user to create a hierarchical visual model representing a business of interest, said model including entities within said business, attributes assigned to said entities and relationships between said entities; and

    a table generator to create database tables for said backend database, said database tables following the hierarchy of said visual model.

26. A modelling tool as defined in claim 25 wherein said table generator creates a table object for each of said entities in said model, said table generator adding database field definitions to each of said table objects corresponding to attributes assigned to the associated entity and corresponding to the relationship the associated entity has with another entity, said table objects being used to create said database tables.

27. A modelling tool as defined in claim 26 wherein said table generator generates a request for a driver of said backend database to generate said database tables using said table objects.

28. A modelling tool as defined in claim 27 wherein said table generator adds database field definitions to said each of said table objects to provide distributed database support.

29. A modelling tool as defined in claim 28 wherein said database field definitions include primary key, last arrived date and last altered date fields.

30. A modelling tool as defined in claim 27 wherein said user interface includes a database selection window to enable a backend database to be selected from a list of available backend databases.

31.   A modelling tool as defined in claim 26 wherein said relationships between entities include "one to many" and "is derived from" relationships, said user interface including user-selectable icons corresponding to the "one to many" and "is derived from" relationships to enable relationships between entities in said model to be defined.

32.   A modelling tool as defined in claim 31 wherein said user interface further includes user-selectable icons to allow defined relationships between entities to be changed.

33.   A modelling tool as defined in claim 32 wherein said visual model is presented on said model creation palette as an entity-relationship diagram, the visual appearance of connections between entities denoting the relationships therebetween.

34.   A modelling tool as defined in claim 33 wherein said user interface includes user-selectable controls to group entities within said model into categories.

35.   A modelling tool as defined in claim 34 wherein said user interface includes user-selectable controls to adjust the appearance of entities in individual groups.

36.   A modelling tool as defined in claim 35 wherein said user interface includes a prefix tool to allow entities in individual groups to be assigned a common prefix.

37.   A modelling tool as defined in claim 26 wherein attributes assigned to said entities are selected from the group consisting of data attributes, calculation attributes and lookup attributes.

38.   A document creation tool for creating a document including at least one data control having a field to receive data, said document creation tool comprising:

a user interface including a document creation palette presenting a document to be created and user-selectable controls to enable a user to select data control types associated attributes from a data schema corresponding to a hierarchical model of a business entity and to create data controls on said document; and

5   a linking engine creating objects to live-link data controls created on said document to database tables of an underlying database corresponding to selected attributes of said data schema.

39.   A document creation tool as defined in claim 38 wherein said linking
10   engine automatically creates said objects when a selected attribute is placed on said document and an associated data control has been created.

40.   A document creation tool as defined in claim 39 wherein said user interface includes a window to present said data schema, selected attributes placed on
15   said document being dragged from said window and dropped onto said document using a computer pointing device.

41.   A document creation tool as defined in claim 40 wherein said user interface includes a toolbar having a plurality of icons, each of said icons representing
20   a different type of data control.

42.   A document creation tool as defined in claim 39 wherein said user interface further includes user-selectable tools to partition said document into regions.

25   43.   A method of licensing database management software to a licensee, said database management software including documents having data controls associated with attributes assigned to entities within a model representing the business of said licensee, data to be managed being entered into a database through said data controls, said method comprising the step of:
30   charging a fee to said licensee for said software based on the number of attributes assigned to entities within said model.

44.   The method of claim 43 wherein said fee is also based on the types of
data controls in said documents.

45.   The method of claim 44 wherein said fee is further based on the
5   number of documents provided to said licensees.

46.   A method of licensing database management software to a licensee,
said document-based database management system including documents created for
said licensee, each document including at least one data control having an associated
10   field into which data is entered, said at least one data control being linked to a
database and populating said database with data entered into said field, said data
controls being associated with attributes assigned to entities within a model of the
business of said licensee, said method comprising the step of:
          charging a fee to said licensee based on the number of attributes
15   assigned to entities within said model.

47.   The method of claim 46 wherein said fee is further based on the
number of documents provided to said licensee.

20   48.   A method for managing data stored in a relational database comprising
the steps of:
          developing a hierarchical model of a business schema;
          translating the model into hierarchical database tables for said
database, said database tables modelling said business schema;
25   creating documents including data controls having fields into which
data is to be entered using the hierarchy of said database tables; and
               live-linking the data controls in said documents to the database tables
so that data entered into said fields automatically populates the database tables.

30   49.   The method of claim 48 wherein said documents are used across a
distributed data network and wherein records created in databases across said data
network are reconciled in a central database.

-54-

50.   The method of claim 49 wherein said business schema follows an oil company and wherein documents include data controls to receive data relating to the business of said oil company.

5

51.   The method of claim 50 wherein documents include data controls to receive data relating to oil wells from startup to completion.

52.   A method of transmitting data to a remote destination over a

10   telecommunications link comprising the steps of:

breaking a data file to be transmitted into packets;

transmitting each packet to said remote site and for each transmitted packet, generating an entry corresponding thereto;

monitoring the amount of data in said file that has been transmitted to

15   determine a current transmission point within said data file;

when said remote site acknowledges receipt of a transmitted packet, deleting the entry corresponding to said packet; and

if said telecommunications link terminates before said data file is fully transmitted, re-establishing said telecommunications link, re-transmitting packets

20   associated with remaining entries and then resuming transmission of packets beginning from the current transmission point at the time the telecommunications link was terminated.

53.   The method of claim 52 wherein said generated entries are maintained

25   in a table and wherein entries in said table are deleted when said destination acknowledges receipt of the packets corresponding thereto.

54.   A file transmission application for transmitting data to a remote site over a telecommunications link comprising:

30   communications to establish a telecommunications link with said remote site;

a packetizer for breaking a data file to be transmitted into packets; and

-55-

a packet handler for transmitting each packet to said remote site, said packet handler generating an entry corresponding to each transmitted packet and monitoring the amount of data in said file that has been transmitted to determine a current transmission point within said data file, said packet handler further being

5 responsive to packet acknowledgments received from said remote site confirming receipt of packets and deleting entries corresponding to acknowledged packets, said packet handler re-transmitting packets associated with remaining entries and resuming transmission of packets beginning from the current transmission point when the communications re-establishes a telecommunications link that was terminated before

10 the data file was fully transmitted.

100

102

DEVELOPMENT TOOLS

| PACKAGER | MODELLER | WORKVIEW TOOLS | DOCUMENT TOOLS | NOT SHIPPED WITH PRODUCT |

106   104   108   110

122

| TELECOMM LINK | DATABASE PACKAGES | INTELLIGENT MODEL | WORKVIEW INTERFACE | DOCUMENTS | OLE |

118   115   111   124   126

LINKING ENGINE

EXTERNAL APPLICATION

116

112

WORD EXCEL

120

114

DATABASE(S)

| ODBC | ACCESS | PARADOX | SQL | ORACLE | INTERBASE | PAS |

Fig. #1

X-CORE
MODEL
DESIGN
TOOL KIT
140

CREATES →

VISUAL DATA MODEL
142

→

CLICK ON A BUTTON TO TRANSLATE VISUAL DATA MODEL TO DATABASE STRUCTURE
144

→

ORACLE INTERBASE SYSBASE PARADOX ACCESS
114

VISUAL DATA MODEL
150

FEEDS →

X - CORE APPLICATION ENGINE
152

→

DRAG AND DROP DATA FIELDS FROM VISUAL DATA MODEL TO CREATE DOCUMENTS
154

→

DOCUMENTS
126

ENTER DATA DIRECTLY INTO DOCUMENTS
160

→

DOCUMENTS
162

→

X - CORE APPLICATION ENGINE
164

→

ORACLE INTERBASE SYSBASE PARADOX ACCESS
114

Fig. 2

Fig. 3

Fig. 4

Fig. 5

Fig.6

Fig. 7

230

Fig. 8a

Fig. 8b

Fig. 8c

Fig. 9a

250

252

260

Fig. 9b

Model Builder 1.0.11.82 - Sample
File Edit View Special

CORPORATION

OBST_CODE

DIVISION

Edit Attribute

Actual Name : DIFFERENCE
Kind : Floating-point
Link Group : (None)
Seg
Init

Display Name : Difference
Size
Decimal Places : 0
Entity

Add spcifier

Script

Cancel

Ok

DIVISION

Add
Add Calc
Add Lookup
Delete
Close

Primary Key
Last Altered Date
Last Arrived Date
Division Name
Division Budget
Division Expenditures
Difference
Key to CORPORATION

Script Editor

Text Location
File
Internal

File Name :

```
function Perform: Real
var
  Link : TEntityLink
endvar
  Link.Get('ADIVISION')
If Link.Synchronize
  Perform := Link.FloatField('DIVISION_ACTUAL')-Link.FloatField('
else
  Perform := 0.0
endif
endproc
(----)
procedure Setup
var
  Interim : Pointer
```

Browse
Save
Save As
View Entities
Ok

180

Fig. 9c

Fig. 10a

272

Fig. 10b

274

Fig. 11

Model Builder 1.0.0.82 - Sample

File Edit View Special

CORPORATION

COST CODE

DIVISION

OILFIELD

OFFSHORE_OILFIELD

GEN_COST

Prefix Entities

What text do you want to prefix the entity names with?

OK   Cancel

276

$\underline{Fig.12}$

Model Builder 1.0.0.82 - Sample

File Edit View Special

CORPORATION

Unit Manager

File

New Family   Delete   New Unit   Delete   New Group   Delete

282a   282b   284a   284b   286a   286b   OFFSHORE_OILFIELD

282   284   286

280

$\underline{Fig.13a}$

CORPORATION

Unit Manager

File

New Family	Delete		New Unit	Delete		New Group	Delete
Length			feet (ft)			Well Depth	
			metre (m)				

280

288

Group Options

Group Name: Well Depth

Stored:
metre (m)

OK

288a

Display:
feet (ft)

Cancel

288c

288b

Fig. 13b

Fig. 14

290

Select Entity
From Model — 200

Create Table Object
For Entity — 202

Create Field Definitions
For Entity Attributes — 204

Add Field Definitions
For Entity Master
Relationships — 206

Add Field Definitions
For Distributed
Database Support — 208

Add Index Definitions
For Master and Primary
Keys — 210

Ask Database Driver to
Create Database Table
Using Table Object — 212

Y — More
Entries? — 214

N

End

Fig. 15

14

Fig. 16

330

332a 332e 332c 332g 336a 336b 336c

Toolbar ☒

332b 332d 332f 336d 336f

332h 336e 336g

Fig. 17

Fig. 18a

Fig. 18b

Division Name : Division A

Division Budget : 500000

Division Expenditures : 475001

Difference : -24999

Oilfield Name : Oilfield A

Corporate Cost Codes!

Cost Name	Cost Account
Drilling	100
Completions	101
Facilities	102
Abandonment	103

Cost Name	Amount
Completions	150.00
Abandonment	149.00
Production	1.00
Construction	101.00

View Extract (F3)

Calcs (Now On)

Schema

- CORPORATION
  - Corporation Name
  - COST_CODE
    - Cost Name
    - Cost Account
  - DIVISION
    - Division Name
    - Division Budget
    - Division Expenditures
    - Difference
  - OILFELD
    - Oilfield Name
    - OFFSHORE_OILFIELD
    - Wave Height
    - Revenue #01
      - GEN_COST
        - Cost Item
        - Cost Account
        - Amount
      - OFFSHORE_OILFIELD
        - Wind
        - Wave Height

OFFSHORE_OILFELD.WAVE_HEIGHT

Sample Workbook
- Corporate Reports
  - Sample Document
  - Sample Document 2

Fig. 18c

Division Name Division A

Division Budget: 500000

Division Expenditures: 475001

Difference: -24999

Oilfield Name Oilfield A

Corporate Cost Codes!

Cost Name	Cost Account
▶ Drilling	100
Completions	101
Facilities	102
Abandonment	103

Cost Name	Amount
Completions	150.00
▶ Abandonment	149.00
Production	1.00
Construction	101.00

Sample Workbook
Corporate Reports
Sample Document
Sample Document 2

Toolbar

New Band

Tools View Entries (F3)

Calcs (Now On)

Schema

CORPORATION
Corporation Name
COST_CODE
Cost Name
Cost Account
DIVISION
Division Name
Division Budget
Division Expenditures
Difference
OILFIELD
Oilfield Name
OFFSHORE_OILFIELD
Wind
GEN_COST
Cost Item
Cost Name
Cost Account
Amount
OFFSHORE_OILFIELD
Wind
Wave Height

OFFSHORE_OILFIELD.WAVE_HEIGHT

Sample Document (C:\Sample\SampleDocument.wcd)

Sample Workbook
Corporate Reports
Sample Document
Sample Document 2

Division Name :Division A

Division Budget : 500000

Division Expenditures : 475001

Difference : -24999

Oilfield Name :

Corporate Cost Codes!

Cost Name	Cost Account
Drilling	100
Completions	101
Facilities	102
Abandonment	103

Cost Name	Amount
Completions	150.00
Abandonment	149.00
Production	1.00
Construction	101.00

Wave Height

350

View Entries [F3]

Calcs (Now On)

Schema
CORPORATION
  Corporation Name
  COST_CODE
    Cost Name
    Cost Account
  DIVISION
    Division Name
    Division Budget
    Division Expenditures
    Difference
  OILFIELD
    Oilfield Name
    OFFSHORE_OILFIELD
      Wind
      GEN_COST
        Cost Item
        Cost Name
        Cost Account
      OFFSHORE_OILFIELD
        Amount
        Wind
        Wave Height
  OFFSHORE_OILFIELD.WAVE_HEIGHT

Fig. 18d

Sample Document [C:\Sample\SampleDocument.wcd]

Welcome 1 0 0 05 (User licence 99 - 0.0% used)

File Edit Model View Options Query Bookmarks Workbook Special

Sample Workbook
Corporate Reports
Sample Document
Sample Document 2
Sample Document 3
Sample Document 4

Maple Leaf Resources

HEADER BAND

Division Name : Division A

[Header]

BODY BAND

Oilfield Name : Oilfield A

[Body]

FOOTER BAND

Division Budget : 500000

Division Expenditures : 475001

[Footer]

New Docd

Sample Document 3 [C:\Sample\Sample Document 3.wcd]

Fig. 19a

WellZONE 1.0.0.85 (User licence 99 - 0.0% used)

File Edit Model View Options Query Bookmarks Workbook Special

Sample Workbook
Corporate Reports
Sample Document
Sample Document 2
Sample Document 3
Sample Document 4

Maple Leaf Resources

HEADER BAND

Division Name : Division A   [Header]

BODY STRIP BAND

Oilfield Name : Oilfield A   [Body]

BODY STRIP CYCLING AREA

New Hard

FOOTER BAND

Division Budget :   500000

Division Expenditures :   475001

[Footer]

Sample Document 4 (C:\SampleView_Sample Document 4.wcd)

Fig. 19b

Fig. 20

Determine Type of Data Aware Control and Create Control Object — 408

Does Entity link Owner For Entity link Object Exist ? — 410

Y

Determine Entity link Object — 411

N

Create Entity link Object — 412

Create Entity Reference Object — 413

Create Attribute Reference link — 414

To Block 416

Fig. 21a

From 414

Hook Data Source Object From Entity Link Object To Control Object — 416

Create Query If Necessary — 418

Generate Field Objects For Query — 420

Poll Entity Link Owner Object To Determine Master Entity — 422

To 424

Fig. 21b

If Master Entity Exists,
Generate Master Link To
Entity Link Object
Corresponding To Master
Entity — 424

Repeat Steps — 426
For all Master Entities

Generate SQL code — 428

Open query — 430

Open Data Source Object — 432
To Live-Link Control

Fig. 21c

Fig. 22a

Fig. 22b

Fig. 22c

Fig. 22d

Bit and Pump Details

Bit

Pump Details

Fig. 23c

Fig. 23b

Fig. 23c

Fig. 23d

Fig. 23e

Fig. 23f

Fig. 24a

Fig. 24b

Fig. 24c

Fig. 25a

Fig. 25b

Fig. 26

Fig. 27a

Fig. 27b

Fig. 27c

Fig. 27d

Down Hole Schematic

Project Type : Sample Wells
Well Name : Sample Well #1 Field Name 1-2-3-4W5
Well ID : 100.1 2 3-4-4 W5M1

As Of Date : 04/16/99

Fig. 27e

Fig. 27f

Fig. 28

Fig. 29a

Fig. 29b

Fig. 29c

Fig. 30a

Fig. 30b

Fig. 30c

Fig. 31a

Drilling Project Cost Variance

Fig. 31b

Fig.32a

Fig. 32b

Fig. 32c

Fig. 33

Fig. 34

Figure 35